# Status of CDF dCache-based Analysis Diskpool

**Grid & Data Management**
**February 28, 2006**

***Doug Benjamin - Duke***

***Dave Ambrose - Fermilab***

# dCache-based Analysis Diskpool: Overview

- **present situation**
  - ⇒ **120 TB of disk space managed by physics groups for ongoing analyses**
    - ⋆ **short-term staging for collection/concatenation/validation**
    - ⋆ **storage of analysis-specific files**
  - ⇒ **large number of file servers with static disk areas**
  - ⇒ **data cataloged on web pages maintained by physics groups**
  - ⇒ **disks accessed via specialized version of rootd**
- **proposed solution:**
  - ⇒ **replace majority of "static project disk" with dCache-based diskpool**
    - ⋆ **product support by Fermilab**
    - ⋆ **global namespace simplifies cataloging, monitoring**
    - ⋆ **maintenance, resource re-allocation w/o distruption**
    - ⋆ **scalable and configurable file serving capability**
    - ⋆ **client analysis software already designed to use for system**

# Current Hardware Setup

- **administrative nodes**

| node | CPU's | RAM | comment |
|---|---|---|---|
| fcdfrdc2 | 4×3066 MHz | 4 GB | pnfs server |
| fcdfrdc3 | 4×3066 MHz | 4 GB | door hosting server |
| fcdfrdc4 | 4×3066 MHz | 4 GB | main dCache server |
| fcdfrdc5 | 4×2666 MHz | 2 GB | monitoring node |

- **pool nodes**

| node | CPU's | RAM | disk space |
|---|---|---|---|
| fcdfdata012 | 2×1393 MHz | 2 Gb | 2 TB |
| fcdfdata034 | 2×1396 MHz | 2 Gb | 2 TB |
| fcdfdata039 | 2×1393 MHz | 2 Gb | 2 TB |
| fcdfdata040 | 2×1393 MHz | 2 Gb | 2 TB |
| fcdfdata119 | 4×3066 MHz | 4 Gb | 7 TB |
| fcdfdata126 | 4×3065 MHz | 4 Gb | 8 TB |
| fcdfdata140 | 4×3202 MHz | 4 Gb | 15 TB |
| fcdfdata141 | 4×3202 MHz | 4 Gb | 15 TB |
| | | | 53 TB |

# Current Usage

- file namespace partitioned into 7 dedicated pool groups

  ⇒ 6 analysis groups plus 1 test pool

| group | assigned | available | used |
|---|---|---|---|
| *stn* | 20 TB | 3.4 TB | 83% |
| *ewk* | 7 TB | 2.5 TB | 65% |
| *top* | 5 TB | 3.2 TB | 35% |
| *qcd* | 5 TB | 1.5 TB | 73% |
| *bnt* | 1.6 TB | 1.4 TB | 12% |
| *exo* | 1.8 TB | 1.8 TB | 0% |

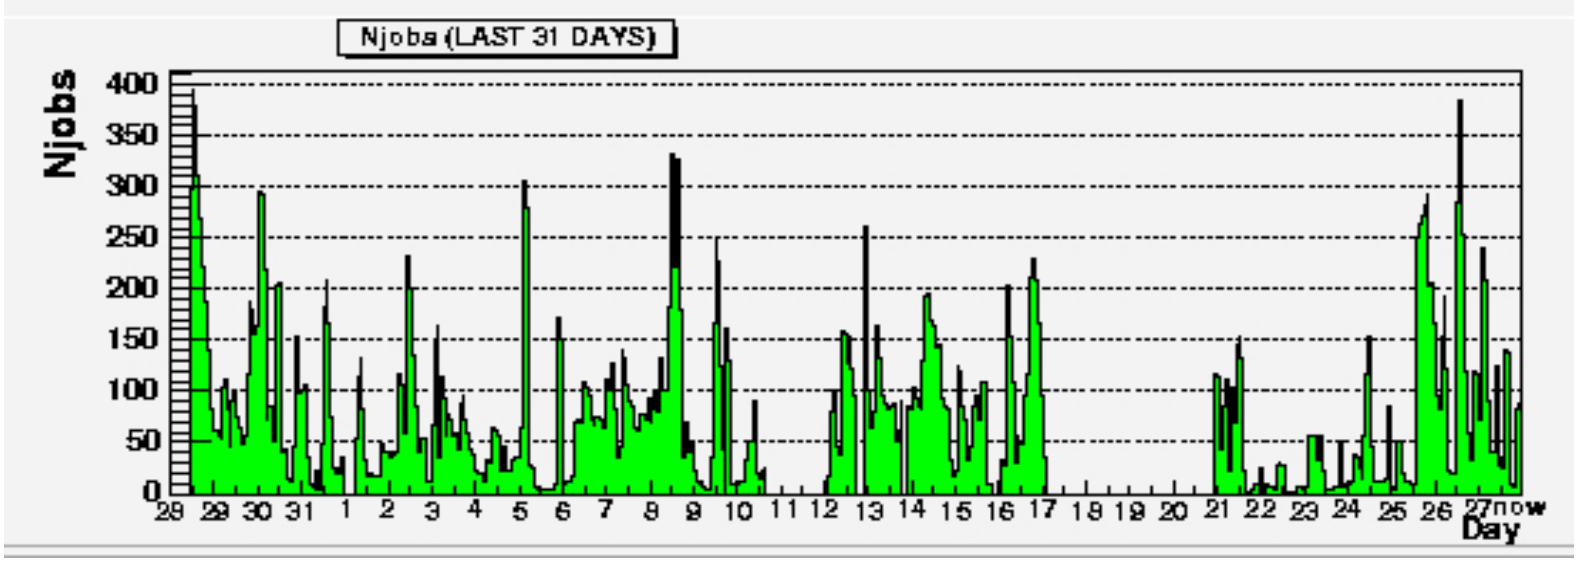

# Current Usage

- **TB read per day (Jan 27 to Feb 27, 2006)**

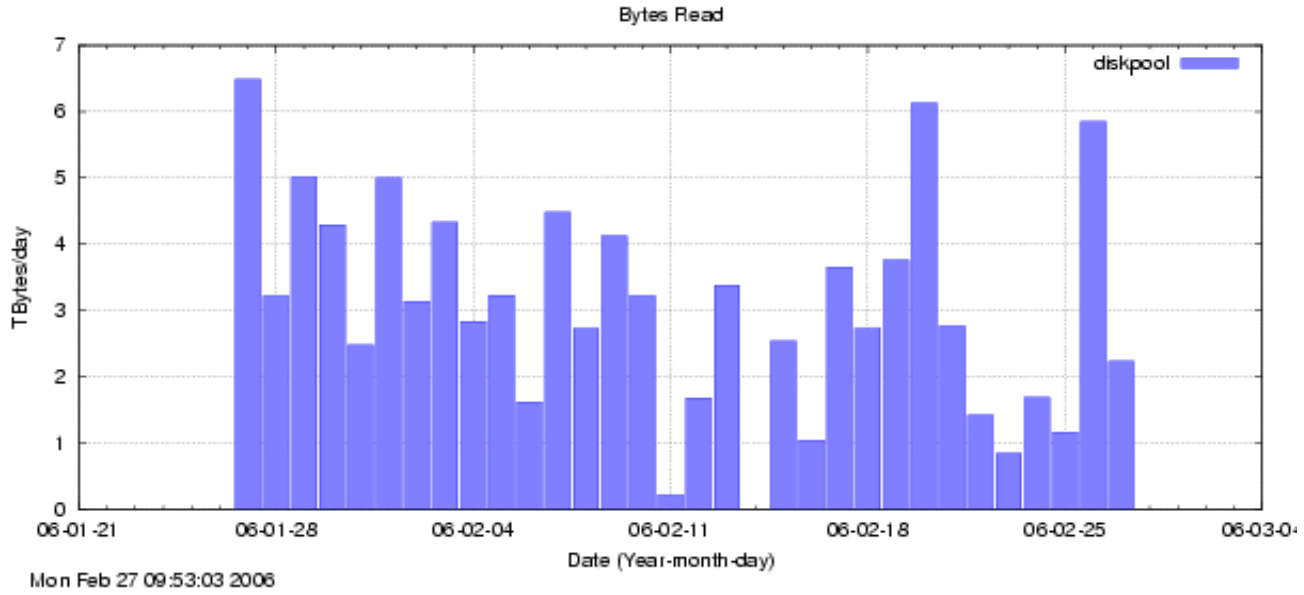


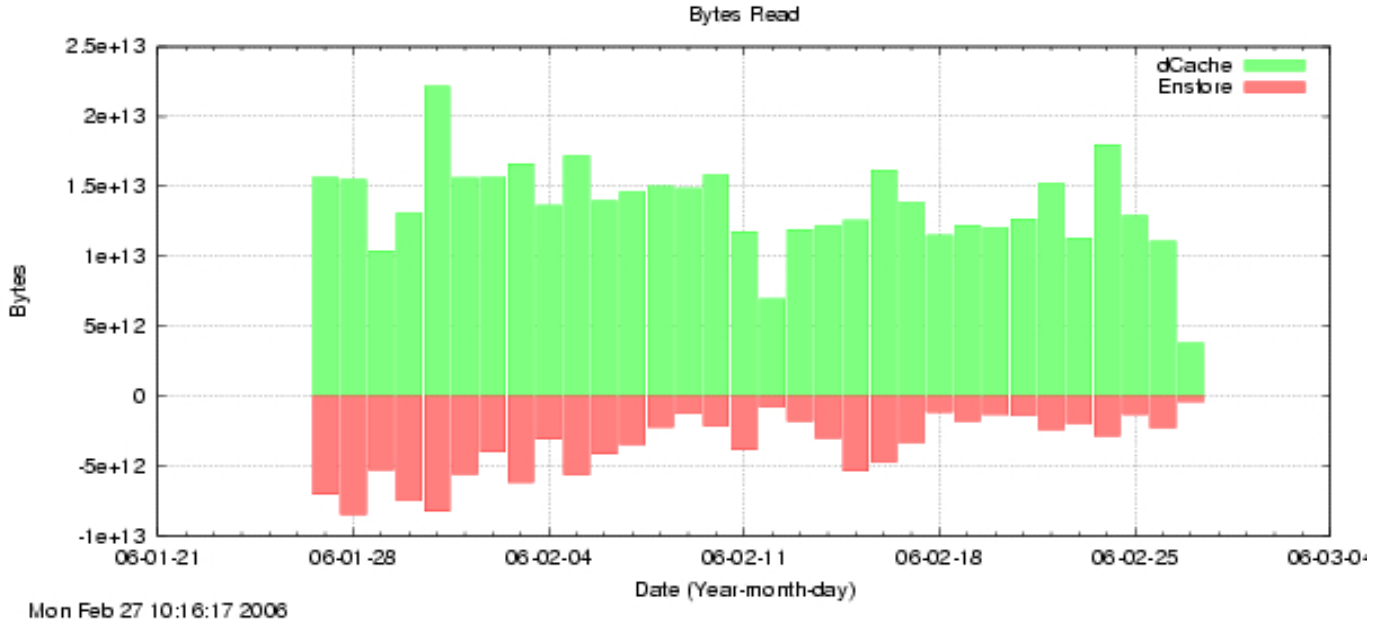

# Status of Deployment

- **deployment plan**
  - ⇒ **Phase 1: prototyping, testing, and understanding performance**
  - ⇒ **Phase 2: pre-production**
  - ⇒ **Phase 3: production operation**
- **Phase 1 goals for 50 TB system**
  - ⇒ **understand characteristics/performance of system**
    - ⋆ **test functionality, load, scalability; investigate failure modes**
    - ⋆ **develop automatic monitoring (based on production dCache system)**
  - ⇒ **understand use cases and system requirements**
    - ⋆ **recruit and train power users**
    - ⋆ **develop usage rules (minimum filesize, prohibit tar, unzip, ls, ...)**
    - ⋆ **estimate how system load will scale with additional users**
  - ⇒ **develop specs for production hardware for Phase 2 ($\sim$100 TB)**
  - ⇒ **develop support agreements for Phase 2**
    - ⋆ **primarily from CDF collaboration, expert help from CD**
  - ⇒ **hold pre-production readiness review**